



Direct or Transparent Proxy?

Choose the right configuration for your gateway.

Table of Contents

Direct Proxy	3
Transparent Proxy	4
Other Considerations: Managing authentication made easier	4
SSL scanning encrypts web traffic	5
Implement protocol scanning for applications that use non-standard ports	5
How proxies perform domain name service lookups	6
Protect off-network mobile PCs	6
Comparison	6
Recommendations	7

The Internet is an essential part of most company's business infrastructure. However, it can be a risky place, but there are ways to minimize risks so your business can thrive. Analysts encourage organizations to deploy a secure web gateway (SWG) to protect their networks against access by malicious software.

Initially, the Internet was designed with the assumption that connections between entities are *explicit* and *stateful*. An SWG intercepts and examines inbound and outbound web traffic and, in effect, becomes a proxy for the user, who no longer interacts directly with the web site. At first, browsers had to be explicitly configured to use the web gateway, which led to the concept of a direct proxy.¹

As networks grew and endpoint management became increasingly complex, the need emerged to control web communications without directly manipulating the endpoint. This led to the concept of a transparent proxy.

The key difference between a direct (or explicit) proxy and a transparent proxy is that a direct proxy is known to the application, which realizes it is talking to a proxy and not the destination server, whereas transparent proxy mode is an intercept model and requires fewer changes to be implemented on the endpoint. Applications think they are going straight to the destination but, in reality, a network service is redirecting the traffic to the proxy, which then forwards it to its destination.

Regardless of the chosen mode, they both require some effort to work within the context of the original design of the web.

Direct Proxy

The use of a direct proxy clearly conforms to the relevant Internet standards. The data flows are well defined, browser features work out of the box, and deployment is relatively easy and straightforward. All connections are stateful, with point-to-point connections, simplifying troubleshooting.

The requirement, however, is that endpoint desktop browsers need to be configured to use the proxy. One way is to manually modify the browser proxy settings. In most standard browsers, you can specify the proxy address and port. Many browsers also let you specify separate ports for different protocols (FTP, HTTP, HTTPS, and others).

A large organization with thousands of desktops, however, may prefer to automate the browser configuration process. One way is to use the web proxy auto-discovery (WPAD) protocol. This requires the deployment of a proxy auto-configuration (PAC) file, which needs to be written and maintained, on the endpoint or the WPAD host server.

Another common tool is to use the group policy object (GPO) editor (which requires Microsoft Active Directory) and limit it to Microsoft Windows-based browsers that leverage Windows Internet settings. GPO can either identify the location of a PAC file for browsers within your network or enforce and configure the manual proxy settings in each browser within your network. Internet Explorer and Chrome both support GPO. Firefox users will need a specially built version of Firefox and a GPO extension.

Direct proxy settings are not the exclusive domain of browsers. Skype and other applications may not support browser-based direct proxy settings by default. However, while most applications support the use of direct proxies, an application that is not proxy-aware, a piece of malware, or a user with a little bit of network expertise can bypass a direct proxy if they are off the corporate network or if the corporate firewalls are not secured correctly.

Pros	Cons
<ul style="list-style-type: none">• Based on RFC standards, supports defined, stateful connections.• Supported by all operating systems and browsers.• Most security features work out of the box.• External DNS resolution not required at the client.• Authentication is straightforward.• Standard troubleshooting and workflows.	<ul style="list-style-type: none">• Desktop setup required.• Some applications, in addition to browsers, must have proxy settings explicitly set.• Can potentially be bypassed by applications that don't support proxy settings or knowledgeable users.

Table 1. Pros and Cons of Direct Proxy

Transparent Proxy

The biggest advantage of a transparent proxy is that it is network-based and seamless from a desktop's perspective. There is a reduced requirement to modify multiple desktops to configure the proxy (see below for exceptions). Applications reach out to their Internet destination as though there was no intervening device.

There are several methods for configuring a transparent proxy. The most popular option is to use web cache communication protocol (WCCP), which is designed to transparently redirect network traffic. Another one to use is some form of Layer 2 redirect (often referred to as policy-based routing), which will redirect traffic at a firewall, router, or switch to a proxy. Finally, there are network modes, such as bridge and router mode. However, in those modes, the proxy is both an infrastructure device and a network security device.

WCCP is the most common way to implement transparent redirection. It is well documented, uses a standardized protocol, and doesn't require any specific infrastructure features, such as routing or bridging.

Pros	Cons
<ul style="list-style-type: none">• Seamless from a desktop perspective.• No desktop setup required.	<ul style="list-style-type: none">• Adds complexity from a network dataflow perspective.• Security filters require workarounds to work.• Authentication is more complex and less robust and accurate.• Protocol selection is key to security.• Complex to troubleshoot and to manage.

Table 2. Pros and Cons of Transparent Proxy

Additional Considerations: Managing Authentication Made Easier

Managing authentication in a transparent proxy environment is somewhat more complex than with a direct proxy.

If you don't want to require users to enter their Windows credentials every time they open a new browser window, you may need to add the proxy IP/hostname into the trusted sites list of your browser. The web gateway itself must be trusted, as the gateway redirects the user to an authentication server hosted on the web gateway itself, which then asks the browser for credentials. The browser requires trust to pass the credentials without user interaction. Failure to trust the SWG IP/hostname will force the browser to prompt the user to enter their credentials for every session.

However, authentication with prompting (using an authentication server) can be accomplished (for Internet Explorer and Chrome) in a transparent deployment with no endpoint configuration changes. The trick is to have a domain name server (DNS) that resolves a host name (short name),

uses the short name in the redirect, and places the SWG IP on the same subnet as the client. With this configuration, the gateway will be a trusted part of the Intranet zone and credentials will be supplied following your default settings.

Authentication needs to be “session-based” in a transparent deployment. In one model (time/IP-based session), the browser has a session bound to its IP address for a certain amount of time. This can cause issues in environments where you either have shared workstations (user A's session is still valid when user B starts using the workstation) or where you have multiple users with the same IP address (Citrix/terminal servers, for example). In these environments, you will have to rely on cookie authentication.

With cookie authentication, you use a cookie that is only valid for the life of the browser session. This requires you to enable the browser to receive third-party cookies—otherwise, users will likely see malformed and incomplete web pages.

SSL scanning encrypts web traffic.

Most modern web applications use the secure socket layer (SSL) protocol to encrypt web (HTTPS) traffic. If you don't implement SSL scanning, you risk allowing malicious software to slip into your network unseen and you lose the ability to control the web application and content that is being sent or received.

To implement SSL scanning on the SWG, you will need to import the SWG root certificate authority (CA) into the browser as a trusted certificate authority. Otherwise, every time a user browses to an SSL-enabled site, they will be prompted to trust the certificate. This applies to both direct and transparent deployments.

In a transparent configuration, when a request for an HTTPS site comes in, it will come through as an IP address instead of the domain name. Some browsers and applications implement server name indication (SNI), but this is not required. If SNI is not present, the SWG will use that IP address as the domain name, which can cause problems when trying to whitelist or block a domain name, or when searching for the domain name in an access log. To get the correct domain name (without SNI), an administrator can add a rule that takes the common name from the certificate and uses it as the domain name, so the whitelist or blacklist works correctly.

Implement protocol scanning for applications that use non-standard ports.

Determining which protocols and ports to scan is important since an application may avoid traffic scanning and filtering by not using a standard web port.

For example, Skype, a popular Voice-over-Internet Protocol (VoIP) service, hops among ports, does not conform to HTTP/HTTPS protocols, and uses peer-to-peer connections. Since Skype can be used to transmit or receive data, it may create a potential malware risk that the SWG cannot detect. To allow but still control Skype, based on authentication, you have to use direct proxy or McAfee® Client Proxy (see below).

If the direct proxy approach is taken, McAfee recommends that you set up a separate proxy port on the SWG for this purpose, use authentication on that port, and block (continue to scan) all categorized sites, all sites not referenced by IP, and all high- and medium-risk reputation sites, other than those that are tunneled or allowed specifically for Skype.

Transparent proxies often allow retention of the endpoint IP address (also known as IP spoofing), in which case routing can be problematic and TCP connections are no longer point to point, but transparently intercepted. If problems arise, this scenario can create troubleshooting difficulties.

How proxies perform domain name service lookups.

In a direct proxy configuration, the endpoint client doesn't need to look up a network address using domain name service (DNS). It relies on the SWG to perform any necessary DNS requests to identify the target web server.

In a transparent proxy configuration, DNS lookup is done on the endpoint client, as it is the client machine that connects to the IP address of the web server. When the request makes it to the SWG, the SWG will also do a DNS request before making its request out to the web server. In some situations, this can lead to complications. For example, if the client and SWG use different DNS servers, they may get different results. It can get even worse if the endpoint can resolve a particular host name via DNS, but, for some reason, the SWG cannot. It is strongly recommended that you make sure that your endpoint clients and SWG use the same DNS server(s).

Protect off-network mobile PCs.

Mobile users present a major proxy concern for administrators. It's relatively easy to protect a users' PC when they're connected to a corporate network. However, mobile (traveling) or remote (home-based) users present a particular challenge—how do you force their web traffic to go through the proxy so it can be scanned for malware?

The answer is to either require users to use a VPN and log onto the corporate network before they access any websites from a remote location (such as a hotel Wi-Fi network) or to use a client proxy, such as McAfee Client Proxy. Since it's difficult to force remote users to use a VPN before they access the web, using McAfee Client Proxy, a tamperproof tool that can be easily installed (through McAfee® ePolicy Orchestrator® software [McAfee ePO™]) on a laptop, is preferable. Whenever the user attempts to connect to the Internet, McAfee Client Proxy intercepts the connection and attempts to reroute it through an on-premises web gateway. If it can't contact a web gateway, then it transparently reroutes the connection through cloud-based McAfee SaaS Web Protection. This ensures that the PC is protected against malware infection just as it would be if it were on an internal network. A third option would be to use only McAfee SaaS Web Protection and require all web traffic to go through the SaaS gateway.

Comparison

Functionality	Direct Proxy	Transparent Proxy
Traffic redirection	Proxy settings need to be pushed to each endpoint.	No settings on the endpoint needed, unless you use McAfee Client Proxy or a tamper-proof client for redirection.
Traffic exceptions	Easily added to proxy settings and on the gateway.	Has to be done on the network/intercepting device (or McAfee Client Proxy policy). Changing a proxy versus a router can be done more often with less risk of affecting the business.
Special applications that cannot deal with proxies	Will ignore proxy settings and attempt to go direct.	Might fail when getting intercepted transparently.
Authentication	Native proxy authentication supported by the browser.	If you are not using McAfee Client Proxy, you need to use an authentication server (to redirect endpoint to special authentication URL).
Authentication sessions	No sessions needed, as each TCP connection is authenticated.	If you are not using McAfee Client Proxy, sessions (time-based or via cookies) are needed. Third-party cookies need to be explicitly allowed on each endpoint.
SSL scanner	Need to push a root certificate authority (CA) to each endpoint to avoid browser warnings.	Need to push a root CA to each endpoint to avoid browser warnings.
SSL traffic	Destination hostnames are seen by the filters.	Only destination IPs are seen by the filters, unless the requester uses a server name indication (SNI).
DNS	The endpoint only needs to talk to the SWG, which is responsible for DNS.	The endpoint and the SWG each need to do their own DNS resolution.
Troubleshooting	It's easy to rule out the SWG by temporarily disabling the proxy on the endpoint. More control over the traffic flow (depending which proxy is being used).	To go direct, without the proxy, network changes are required. It's harder to determine which proxy is used by a particular endpoint.

Table 3. Comparing Direct Proxy and Transparent Proxy

Learn More

For more information, we invite you to visit **Direct Proxy versus Transparent Deployment** on the McAfee Community site. For more information about McAfee on-premises and SaaS secure web gateway solutions, visit www.mcafee.com/webprotection.

Recommendations

Based on many years of experience in networks, ranging from small and simple to the largest and most complex, McAfee has some key recommendations for McAfee Web Gateway customers about proxy configuration.

McAfee Web Gateway, at its core, is a proxy product. The most successful deployment method, in our experience, is a combination of direct proxy with an endpoint proxy PAC file. If the network traffic is expected to be heavy or there is a high-availability requirement, McAfee has had good experience using a hardware load balancer in front of two or more McAfee Web Gateway appliances.

McAfee has endpoint management software, such as McAfee Client Proxy and McAfee ePO software which make the process of configuring and managing the endpoint much easier. This gives you enterprise-wide flexibility when it comes to change management, high availability, and load balancing, as well as easier maintenance and troubleshooting.

If you decide to use a transparent proxy method other than McAfee Client Proxy, McAfee recommends using WCCP. WCCP has many enterprise features already built in, and we often see it in addition to the direct proxy deployments. McAfee supports both modes running simultaneously on the same appliance.

If you decide to deploy either transparent router or transparent bridge modes, the most common support issue we see is with exceptions. Everyone needs exceptions at some point (internal resources, non-standard compliant applications, and so on). In transparent router mode, you can “route” around McAfee Web Gateway on a logical level. In transparent bridge mode, McAfee Web Gateway is in the physical path of the traffic (inline), and it is almost impossible to get around it. An important consideration when using bridge mode is the potential for a complete network outage if no precautions are taken before a McAfee Web Gateway appliance goes offline (such as for an upgrade or maintenance).

Here's a summary:

- For controlled Windows clients, use McAfee Client Proxy.
- For non-Windows controlled clients Mac, Linux), use WPAD/PAC/explicit settings.
- For uncontrolled clients, use WCCP or other transparent redirect to McAfee Web Gateway (on premises) or McAfee SaaS Web Protection (cloud).

About McAfee

McAfee is now part of Intel Security. With its Security Connected strategy, innovative approach to hardware-enhanced security, and unique Global Threat Intelligence, Intel Security is intensely focused on developing proactive, proven security solutions and services that protect systems, networks, and mobile devices for business and personal use around the world. Intel Security combines the experience and expertise of McAfee with the innovation and proven performance of Intel to make security an essential ingredient in every architecture and on every computing platform. Intel Security's mission is to give everyone the confidence to live and work safely and securely in the digital world. www.intelsecurity.com.

